

EV961411962

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

FOR

BANDWIDTH-EFFICIENT PROCESSING OF VIDEO IMAGES

Inventor:

Stephen J. Estrop

ATTORNEY'S DOCKET NO. MS1-1679US

1 This application claims the benefit of U.S. Provisional Application No. 60/492,029,
2 filed on August 1, 2003, and incorporated by reference herein in its entirety.
3

4 **TECHNICAL FIELD**

5 This subject matter relates to the rendering of video data in a bandwidth-efficient
6 manner, and, in a more particular implementation, to the de-interlacing of video data in a
7 bandwidth-efficient manner.
8

9 **BACKGROUND**

10 Video image sequences can be represented using two different techniques:
11 progressive and interlaced techniques. With progressive video, all of the picture elements
12 comprising a frame of a video image are sampled at the same time (where a “frame”
13 refers to one video image snapshot within the sequence of video images). With interlaced
14 video, alternate lines of the video image are sampled at alternate moments in time.
15 Capturing the video image using the interlaced video technique results in the creation of
16 two half-frames, referred to as “fields,” which together constitute a single frame in the
17 video sequence. For instance, note Fig. 1. A first field 102 is commonly referred to as a
18 top field or even field. A second field 104 is commonly referred to as a bottom field or
19 odd field. These two fields (102, 104) together constitute a single frame 106 in a video
20 sequence. Video devices in the United States commonly present video sequences at a
21 frame rate of 29.97 frames per second (that is, roughly 30 frames every second).
22 Televisions and camcorders are common devices that manipulate image sequences using
23 the interlaced technique. Computer monitors and some DVD players are common
24 devices that manipulate image sequences using the progressive technique.
25

1 Some applications require that interlaced video image sequences be processed and
2 displayed using a device that is configured to handle only progressive image sequences.
3 These applications therefore require that the image sequence be converted from
4 interlaced format to progressive format. This process is referred to as “de-interlacing.”
5 Fig. 1 illustrates the result of an exemplary de-interlacing process. In this process, the top
6 and bottom fields (102, 104) are reassembled to form a single frame 108 in the video
7 sequence.

8 A number of techniques exist in the art for de-interlacing video image sequences.
9 However, there is room for improvement in many of these techniques, particularly for
10 those kinds of devices that may be particularly sensitive to variations in bandwidth.
11 More specifically, a technique for de-interlacing video image sequences will inherently
12 require a certain amount of bandwidth as information is transferred to and from memory
13 locations. Modern processors are generally available that can handle such a task in an
14 efficient manner. However, in many competitive marketing environments, the success of
15 a video processing device does not merely depend on the processing power of the device,
16 but also depends on the cost of the device. A developer may therefore choose to forego a
17 certain amount of processing power to provide a less expensive design. It is particularly
18 in these circumstances that the ability to efficiently perform de-interlacing may be
19 jeopardized.

20 The effectiveness of the de-interlacing operation is particularly brought into
21 question in those cases where less powerful architectures are required to perform other
22 processing tasks in combination with the de-interlacing operation. For instance, modern
23 video processors often require functionality that allows for the rendering of supplemental
24 data along with the sequence of video images. Such supplemental data may comprise
25 close-captioning information, various graphical data associated with the presentation of

1 DVD video sequences, PAL Teletext images, and so on. Such supplemental information
2 is referred to as “video sub-stream” data herein, to distinguish this data from the main
3 sequence of video images (referred to as the “video stream” data herein). It may strain
4 the processing resources of some architectures to handle the combined tasks of de-
5 interlacing and the rendering of video sub-stream data. Similar potential problems may
6 be present with respect to other kinds of resource-intensive processing tasks, such as the
7 display of high definition television (HDTV) signals.

8 Consider, for example, the case of a device that employs Unified Memory
9 Architecture (UMA). In a UMA design, both the main CPU processor of the device and a
10 graphics processing unit (GPU) share the same memory. (A GPU commonly provides a
11 pipeline for performing a sequence of image rendering tasks, therefore alleviating some
12 of the processing burden that would otherwise be imposed on the CPU). Allocating
13 separate memories to the CPU and GPU will typically result in a more robust processing
14 architecture. However, many developers opt for the UMA architecture so as to provide a
15 less expensive design. This savings in cost can result in a higher bandwidth burden on
16 the device, as both the CPU and GPU are now making demands on the shared memory.
17 It is in these kinds of devices that limitations may arise that prevent video processing
18 operations from becoming too complex. Some UMA-type devices may not be able to
19 efficiently perform de-interlacing and the processing of video sub-stream data at the same
20 time. This can result in the suboptimal display of video data, that is, at less than the
21 normal frame rate. In other cases, these limitations may completely preclude the efficient
22 rendering of video image data.

23 Accordingly, there is an exemplary need in the art to provide techniques for
24 allowing devices to perform video processing without incurring the above-identified
25

1 problems. There is a more specific need in the art to provide techniques for robustly
2 handling video processing in bandwidth-challenged devices, such as UMA-type devices.
3

4 **SUMMARY**

5 According to one exemplary implementation, methods, apparatuses, and computer
6 readable media are described for performing multiple video processing tasks in a single
7 operation, as opposed to serially. For instance, methods, apparatuses, and computer
8 readable media are described for de-interlacing a principal video stream at the same time
9 that at least one video sub-stream is combined with the principal video stream.
10 Performing these tasks in a single call to a data processing module (such as a GPU), as
11 opposed to staggered serial calls, reduces the bandwidth requirements of the processing
12 operation. This, in turn, can enable a computing device to perform these multiple
13 operations at full frame rate.

14 In one implementation, different texturing units are assigned to the principal video
15 stream and the video sub-stream, respectively. The graphics processing unit interacts
16 with these texturing units and associated memory locations substantially in parallel, thus
17 providing the above-described bandwidth savings.

18 The methods, apparatuses, and computer readable media described above can
19 combine other kinds of processing tasks in a single call to the data processing module
20 (e.g., other than de-interlacing). For example, the methods, apparatus, and computer
21 readable media can perform resizing of the principal video stream at the same time as the
22 sub-stream is combined with the principal video stream.
23

24 **BRIEF DESCRIPTION OF THE DRAWINGS**

25 Fig. 1 shows an example of a de-interlacing operation.

1 Fig. 2 shows a graphical processing pipeline that performs de-interlacing in series
2 with combining a video stream with a video sub-stream.

3 Fig. 3 shows a graphical processing pipeline that performs de-interlacing at the
4 same time that it combines a video stream with a video sub-stream.

5 Fig. 4 shows an exemplary apparatus for implementing the graphical processing
6 pipeline shown in Fig. 3.

7 Fig. 5 shows an exemplary allocation of texturing units to components of a video
8 stream and to a sub-stream, respectively.

9 Fig. 6 illustrates an overview of a procedure for performing multiple video
10 processing tasks substantially in parallel.

11 Figs. 7-13 show examples of the operation of a DeinterlaceBltEx method used in
12 the procedure of Fig. 6.

13 Fig. 14 shows an exemplary computing environment for implementing the
14 apparatus shown in Fig. 4.

15 The same numbers are used throughout the disclosure and figures to reference like
16 components and features. Series 100 numbers refer to features originally found in Fig. 1,
17 series 200 numbers refer to features originally found in Fig. 2, series 300 numbers refer
18 to features originally found in Fig. 3, and so on.

19

20 **DETAILED DESCRIPTION**

21 This disclosure pertains to apparatuses, methods, and computer-readable media
22 for processing a video stream. A video stream refers to a primary sequence of video
23 frames. On the other hand, a video sub-stream refers to auxiliary video data that can be
24 presented along with the video frames of the primary video stream. In one case, a device
25 can be configured to automatically present the video sub-stream data in tandem with the

1 primary video stream. In other cases, the device can be configured to render the video
2 sub-stream data only upon the command of a user, or based on some other contingent
3 event. For instance, a television signal may contain interlaced video fields that are
4 supplemented by close-captioning data. This close-captioning data is only displayed
5 when the viewer selects such a display (e.g., via remote control actuation). The close-
6 captioning data is combined with the primary video stream to form a composite image for
7 display.

8 While the ensuing discussing will emphasize the bandwidth-efficient integration
9 of a de-interlacing operation with a video sub-stream compositing operation, the
10 principles described herein have application to other video processing environments and
11 applications, and other combinations of video processing operations. The video
12 processing applications described in the following are exemplary and illustrative.

13 Further, the novel features of the present invention are described below as
14 implemented on a personal computer or like device. However, the principles can also be
15 implemented on other devices, such as game consoles, central server type computers, etc.

16 This disclosure is organized as follows. Section A provides an overview of the
17 video processing techniques that are used to improve the bandwidth-related efficiency of
18 video processing operations. Second B discusses an exemplary apparatus for
19 implementing the strategy described in Section A. Section C provides additional details
20 regarding the manner of operation of the apparatus described in Section B. And Section
21 D discusses an exemplary computer environment for implementing the apparatus
22 described in Section B.

23
24
25

A. Overview of Design Strategy

2 Fig. 3 shows an overview of exemplary design principles that can be employed to
3 improve the bandwidth-related performance of a device that renders video stream data.
4 However, an appreciation of the merits of the design shown in Fig. 3 can best be gleaned
5 by an introductory discussion of a related system 200 shown in Fig. 2. The related
6 system 200 shown in Fig. 2 can be implemented using techniques described in commonly
7 assigned copending application 10/273,505, entitled "Methods and Apparatuses for
8 Facilitating Processing of Interlaced Video Images for Progressive Video Displays,"
9 which names Stephen J. Estrop as sole inventor, and which is incorporated by reference
10 herein in its entirety. Generally, in both Figs. 2 and 3, the ovals represent data that may
11 be stored in respective memory locations, and the boxes represent actions performed on
12 such data. More specifically, the boxes can represent actions that require at least one
13 read to memory to receive input data, and at least one write to memory to provide output
14 data.

More specifically, the system 200 shown in Fig. 2 can be implemented, at least in part, using a graphics module coupled to a personal computer, game console, or other kind of processing device. The graphics module may constitute processing logic permanently coupled to the computer, game console, or other kind of device. Alternatively, the graphics module can constitute processing logic that is detachably coupled to the computer, game console, or other kind of device. In the latter case, the graphics module can constitute a graphics processing unit (GPU) card that can be detachably coupled to the computer, game console, or other kind of device. As well known in the art, a graphics processing unit can implement a series of operations directed to the graphical rendering of information. Such operations can include various tasks associated with a conventional three dimensional graphics pipeline. In the specific case

1 of Fig. 2, only a subset of operations pertinent to de-interlacing and integrating video sub-
2 stream data are shown. More specifically, the circles shown in Fig. 2 represent data at
3 various stages in the pipeline, while the rectangles represent operations performed on the
4 data.

5 To begin with, the system 200 begins by receiving compressed video data 202 and
6 decoding the compressed video data 202 in a video decoder 204. Video decoding can
7 constitute converting the compressed video data 202 from its received form (e.g., MPEG-
8 2 format or some other format) to another format suitable for processing, as well as
9 potentially uncompressing the received video data. Video decoding can also constitute
10 performing motion compensation or iDCT operations. As a result of these operations, the
11 video decoder 204 yields uncompressed video data 206. This video data 206 constitutes
12 a current frame of video data that is fed into the de-interlacing logic 208. The de-
13 interlacing logic 208 typically works by processing a current frame with respect to one or
14 more previous frames of received video data. To this end, Fig. 2 also indicates that
15 previous uncompressed video data 210 and previous uncompressed video data 212 are
16 fed into the de-interlacing logic 208. The dotted line that separates video data 210 from
17 video data 212 indicates that additional samples of uncompressed video data can be fed
18 into the de-interlacing logic 208. The de-interlacing logic 208 itself can employ the well-
19 known bob and weave algorithm, or another other kind of de-interlacing algorithm. The
20 output of the de-interlacing logic 208 constitutes de-interlaced video data 214. In other
21 words, the de-interlaced video 214 constitutes a progressive frame of video data.

22 Video data is commonly represented in conventional video color space, e.g., using
23 conventional Y, U, V color components. The Y component generally represents the
24 brightness of a pixel in the video image frame. The U and V components generally
25 represent the color of the pixel. More specifically, in one exemplary implementation, the

1 Y component is related to the red (R), green (G), and blue (B) components of a pixel in a
2 video frame according to the formula:

3 Given:

4 $K_r = 0.299$

5 $K_b = 0.114$

6 $L = K_r * R + K_b * B + (1 - K_r - K_b) * G$

7 Then:

8 $Y = \text{round}(219 * L / 255) + 16$

9 $U = \text{round}(224 * 0.5 * (B - L) / ((1 - K_b) * 255)) + 128$

10 $V = \text{round}(224 * 0.5 * (R - L) / ((1 - K_r) * 255)) + 128.$

11 The U component of the pixel is represented by $U = B - Y$, and the V component
12 of the pixel is represented by $V = R - Y$. The color space conversion logic 216 converts
13 the Y, U, and V video components to their counterpart R, G, and B components. This
14 yields RGB video data 218. It is conventional to convert video signals into R, G, and B
15 components because graphics processing functionality is traditionally configured to
16 process image data in this format. The color space conversion logic 216 also can
17 combine the color processing operation with a pixel aspect ratio correction process.

18 Next, composition logic 220 combines the RGB video data 218 with video sub-
19 stream data 222 to yield composite video frame data 224. The video sub-stream data 222
20 can constitute closed caption data, any type of supplemental information associated with
21 the presentation of DVD video information, PAL Teletext images, and so on. Although
22 nothing in this disclosure limits the form that the video sub-stream data 222 can take,
23 current manifestations of such data 222 typically take a much simpler form compared to
24 the received primary video data. For instance, this sub-stream data 222 most commonly
25 takes the form of simple graphical overlay data.

Finally, renderer functionality 226 processes the composite video frame data 224 to provide rendered video 228 which is output to a display device (not shown). At this stage in the pipeline, the system 200 can also alter the transparency of each stream and, if required to do so, horizontally or vertically invert the video image. At this stage, the system 200 can also resize the final image and draw it to the implementing device's display. This can be performed by writing data directly to the implementing device's frame buffer or to the device's back buffer which is flipped to the foreground upon a next vertical retrace signal.

The system 200 shown in Fig. 2 works well providing that sufficient processing power is incorporated into the device that implements the system 200. However, as noted in the Background section of this disclosure, developers often opt to reduce the processing power of the implementing device in order to provide a more economical, and thus potentially more marketable, device. For instance, an implementation that devotes separate memory to each of the CPU and GPU used by the implementing device will typically have better performance than an implementation that provides a single memory for shared use by the CPU and GPU (that is, a Unified Memory Architecture design). The operations performed by system 200 may be problematic when performed on such lower-end or medium-end architectures.

More specifically, the design implemented by the system 200 requires at least two significant processing operations performed in series. First, the de-interlacing logic 208 de-interlaces a collection of video samples (e.g., samples 206, 210, 212, etc.). This imposes a first demand on the bandwidth of the implementing device. Then, in a second stage, the composition logic 220 adds the sub-stream video data 222 to the RGB video data 218. This imposes a second demand on the bandwidth of the implementing device. If multiple pieces of video sub-stream data need to be added to the RGB data 218, then

1 the composition logic 220 may further require multiple passes to execute its function. All
2 of this can begin to overload the implementing device. This, in turn, can result in the
3 failure to render the principal video data and video sub-stream data 222 at a normal frame
4 rate. These problems are compounded by the conversion of the YUV data to RGB data.
5 RGB data often takes more bits to represent the color compared to YUV data, thus adding
6 extra overhead to the operations that follow the color space conversion 216.

7 Fig. 3 shows a system 300 that solves the potential problems shown in Fig. 2.
8 Fig. 3 starts with the same operations as the system 200 shown in Fig. 2. Namely, in Fig.
9 3, compressed video data 302 is provided to a video decoder 304 to yield a current frame
10 of uncompressed video data 306. Further, also like the case of Fig. 2, the current frame
11 of uncompressed video data 306 is sent to de-interlacing logic 308, along with one or
12 more previous samples of uncompressed video (e.g., previous uncompressed video data
13 310 and 312). However, unlike the case of Fig. 2, the logic 308 shown in Fig. 3 also
14 functions to add video sub-stream data (e.g., from video sub-stream data 314, 316, etc.) to
15 the video data that is in the process of being de-interlaced. In other words, the logic 308
16 effectively combines the de-interlacing operation with the sub-stream compositing
17 operation. Further, as will be described in greater detail below, the logic 308 executes
18 these two functions such that they can be performed in a single memory read/write
19 operation, rather than in multiple passes. That is, whereas the system 200 shown in Fig. 2
20 requires at least two stages to perform the de-interlacing and the compositing operations
21 (requiring at least two reads from memory), the system 300 shown in Fig. 3 requires only
22 a single stage (e.g., a single memory read/write transaction) (that is, requiring, in one
23 exemplary case, only a single read from memory).

24 As will be discussed in greater detail in later sections of this disclosure, the
25 above-described bandwidth efficiency can be achieved by assigning video data (e.g.,

1 video data 306, 310, 312, etc.) and the video sub-stream data (e.g., 314, 316, etc.) to
2 different respective texturing units employed by the GPU module used in the rendering
3 device. In their most common application, texturing units are assigned to image surfaces
4 to be manipulated in the course of a three-dimensional rendering application. For
5 instance, a “texture” generally refers to an image that is “pasted” onto the surface of a
6 geometric primitive (e.g., a triangle) that then forms a part of a three dimensional scene to
7 be rendered. For instance, the texture of a brick wall can be pasted onto a mesh of
8 triangles, and the texture of a shadow can thereafter be pasted on top of that, and so on.
9 Wolfgang F. Engel, *Direct3D: ShaderX: Vertex and Pixel Shader Tips and Tricks*, 2002,
10 Wordware Publishing, Inc. provides background information regarding this technology.
11 These different texture surfaces are assigned to different so-called texturing units. GPUs
12 are designed to process image data provided by the texturing units in a particularly
13 streamlined and parallel manner to facilitate the rapid display of 3D data, for instance, in
14 a gaming environment which demands real time rendering. The system shown in Fig. 3
15 assigns video data (e.g., 306, 310, 312, etc.) and video sub-stream data (314, 316, etc.) to
16 respective texturing units and thus achieves similar efficient performance of the de-
17 interlacing and compositing operations. More specifically, the GPU logic essentially
18 reads and processes data from each of a collection of texturing units at the same time,
19 instead of in a staggered serial fashion. This reduces the bus congestion in the
20 implementing device, and better enables the implementing device to present the video
21 data and associated video sub-stream data at an optimal frame rate.

22 In other implementations, the logic 308 can provide other processing functions
23 besides de-interlacing and compositing, such as image resizing. In other
24 implementations, the logic 308 can replace one or more of the above-enumerated
25 operations with other kinds of operations. For instance, if the video data is already in

1 progressive format, the logic 308 can be dedicated to simply resizing the video data, or
2 performing some other desired operation on the video data.

3 Continuing with the discussion of Fig. 3, the output of logic 308 includes
4 composite de-interlaced video data 318. This video data 318 is still in YUV color space
5 at this time, as opposed to RGB color space. Color space conversion logic 320 converts
6 the YUV data into RGB data and renders this data to the display to yield rendered video
7 data 322. Note that, compared to Fig. 2, the video data is maintained in YUV form
8 farther along into the video pipeline. This offers additional potential bandwidth savings.
9 For instance, in one implementation, YUV data can be presented using 12 bits (8 bits for
10 the Y component, 2 bits for the U component, and 2 bits for the V component). In
11 contrast, in one implementation, RGB data requires 32 bits to represent. Hence,
12 performing such tasks as compositing using YUV data instead of RGB data cuts down on
13 the bus traffic association with this transaction.

14 Experiments indicate the system 300 shown in Fig. 3 achieves a 62% bandwidth
15 saving compared to the system 200 shown in Fig. 2.

16 **B. Exemplary Apparatus for Implementing the Design Strategy**

17 *Overview of Apparatus*

18 Fig. 4 shows an overview of an exemplary apparatus 400 that can be used to
19 implement the design solution shown in Fig. 3. The apparatus 400 particularly relies on
20 the DirectX® family of technologies produced by Microsoft® Corporation of Redmond,
21 Washington. The DirectX family includes, DirectX Video Acceleration (DirectX-VA),
22 Direct3D, DirectDraw, etc. However, the principles described above can be performed
23 using other kinds of rendering technologies that run on other kinds of operating systems
24 besides the Windows® operating system. The apparatus 400 can represent a personal
25

1 computer, a game console (such as Microsoft® Corporation's Xbox™ gaming console),
2 or other kind of device.

3 To begin with, the apparatus 400 accepts video data from any one of a number of
4 sources. For example, the apparatus 400 can accept video data from a network 402 (such
5 as a remote source coupled to the Internet), any kind of database 404, any kind of
6 computer readable disc media 406 (such as an optical disk, DVD, etc.), or some other
7 source 408. Although not shown, the video data can also be received via wireless
8 broadcast signal or some other broadcast format. The video data is typically stored and
9 transmitted in compressed format (e.g., in one of the MPEG formats, or some other
10 format). However, the received data need not be compressed. The received data will
11 typically contain a combination of video data and audio data. A demux unit 410
12 separates the audio data from the video data. Audio processing functionality 412
13 processes the audio information. However, as this disclosure does not directly pertain to
14 the processing of audio data, further details regarding this functionality are omitted so as
15 not to unduly complicate this disclosure.

16 A video decoder 414 processes the video data. The video decoder can convert the
17 compressed video data from its received format to some other format, as well as perform
18 initial resizing or other operations on this data. The output of the video decoder 414 can
19 include so-called pure video data as well video sub-stream data. The pure video data
20 constitutes the primary video stream to be rendered on the display device. The video sub-
21 stream data can constitute any supplemental information associated with the pure video
22 data, such as close-captioning data, any kind of graphical overlay data (such as various
23 graphical editing controls), various kinds of sub-images presented by DVD players, and
24 so on.

1 In one exemplary implementation, a video mixing renderer (VMR) module 416
2 performs a central role in the processing of the thus received video data. By way of
3 overview, the VMR module 416 interacts with a graphics interface 418 and a display
4 driver 420, which, in turn, controls a graphics module 422. As will be described in
5 greater detail in the next section, this interaction involves probing the capabilities of the
6 graphics module 422. This interaction also involves coordinating the processing of the
7 video data by the graphics interface 418, display driver 420, and graphics module 422. In
8 one implementation, the graphics interface 418 can be implemented using the DirectDraw
9 functionality provided by Microsoft® Corporation's DirectX. DirectDraw serves in this
10 context as a messaging conduit for communicatively coupling the VMR module 416 to
11 the graphics module 422. The graphics module 422 itself can constitute a fixed module
12 within a computer or like device, or it can constitute a detachable unit, such as a graphics
13 card. The vertical chain of functionality represented by the VMR module 416, graphics
14 interfaces 418, display driver 420, and graphics module 422 is divided into a user mode
15 and a kernel mode. As well understood in the operating system art, the user mode refers
16 to aspects of the programming functionality that can be manipulated by the user via
17 various interfaces. The kernel mode represents aspects of the programming functionality
18 that cannot be directly manipulated by the user.

19 The graphics module 422 itself includes one or more graphics processing units
20 (GPUs) 424. A GPU 424 is generally a processing device like the CPU. The GPU 424 is
21 commonly allocated data-intensive rendering tasks that are repetitively performed by the
22 implementing apparatus 400. Allocating these repetitive or data-intensive tasks to the
23 GPU 424 frees the CPU (not shown) to perform other tasks, and thus improves the
24 performance of the apparatus 400. Two exemplary tasks that this invention allocates to
25 the GPU 424 are de-interlacing and rate conversion. These functions are represented by

1 de-interlacing logic 426 and rate converter logic 428. The de-interlacing logic 426
2 combines two or more fields of video data together to form a frame of video data. The
3 rate converter logic 428 modifies the frame rate of a sequence of video frames.

4 The GPU 424 can interact with a local memory 430 associated with the graphics
5 module 422. This local memory 430 can serve any number of storage-related purposes.
6 For instance, this memory 430 can store a final video surface which is then forwarded to
7 a display device 432. The local memory 430 can also store the input surfaces that will be
8 processed by the GPU 422 to yield the final output surface. That is, the local memory
9 430 can be used to implement the texturing feature described above. In addition, or
10 alternatively, the apparatus 400 can rely on shared memory 434 provided by a unified
11 memory architecture (UMA). The unified memory 434 can be shared by the CPU (not
12 shown) and the GPU module 422.

13

14 *Allocation of Texturing Units for Performing Combined Operation*

15 Fig. 5 shows an exemplary organization 500 of texturing units and associated
16 memory. These elements are used to process the principal video data simultaneously
17 within the video sub-stream data. The memory and texturing units can be implemented
18 using the local memory 430 associated with the graphics module 422 shown in Fig. 4,
19 and/or the shared memory 434.

20 The joint operation of interleaving and compositing take place while the video
21 data is still in YUV form. Hence, the VMR module 416 allocates portions of memory for
22 storing this YUV data. A first block 502 can be allocated to storing Y data, a second
23 block 504 can be allocated to storing U data, and a third block 506 is allocated to storing
24 V data. More bits are allocated to the Y data than the U and V components. For
25 instance, for an image that contains 720x480 pixels, a block of 720x480 bytes can be

1 allocated to storing the Y data, a block of 360x240 bytes can be allocated to storing the U
2 data, and a block of 360x240 bytes can be allocated to storing the V data. Finally, a
3 block 508 of memory can be allocated to storing sub-stream data (such as close-
4 captioning data, DVD sub-image display data, graphical icons of various sorts, and so
5 on).

6 In the exemplary organization 500 shown in Fig. 5, only four texturing units are
7 shown (510, 512, 514, and 516). However, other implementations will include more than
8 four units. Texturing unit 510 is allocated to handling an image input surface associated
9 with memory block 502 (i.e., the Y data), texturing unit 512 is allocated to handling an
10 image input surface associated with memory block 504 (i.e., the U data), and texturing
11 unit 514 is allocated to handling an image input surface associated with memory block
12 506 (i.e., the V data). Texturing unit 516 is allocated to handling an image input surface
13 associated with memory block 508 (i.e., the sub-stream data). The memory blocks (Y, U,
14 V, and the sub-stream data) are separate, but need not be contiguous in memory.
15 Additional memory blocks and texturing units can be provided to handle additional video
16 reference samples and/or additional video sub-streams. For instance, an application that
17 includes two previous reference streams would require at least nine texturing units (e.g.,
18 three units for the current sample, and six units for the two reference samples).

19 Finally, Fig. 5 generically shows GPU processing logic 518 associated with the
20 GPU 424 of Fig. 4. The GPU processing logic 518 interacts with the texturing units. The
21 GPU processing logic 518 can perform de-interlacing, frame rate conversion, and/or
22 other task(s).

23
24
25

C. Exemplary Method of Operation

Overview

Fig. 6 shows an overview 600 of the operation of the apparatus 400 shown in Fig.

4. Certain aspects of this overview 600 are common to the processing described in the above-referenced copending U.S. Serial No. 10/273,505. Hence, particular attention will be paid in the ensuing discussion to features that are unique to the joint de-interlacing/compositing functionality of the present invention.

To begin with, in step 602, the VMR module 416 queries the display driver 420 and graphics module 422 regarding what processing modes it supports. After it receives a reply, in step 604, the VMR module 416 sends another inquiry to find out more specific information regarding the capabilities of the display driver 420 and associated graphics module 422. Steps 602 and 604 are discussed in greater detail below under the general heading of “preliminary data processing.”

After investigating the capabilities of the attached hardware and associated interfaces, the VMR module 416 opens a video stream object so that video data and control information can be forwarded to the hardware. Then, in step 608, the VMR module 416 coordinates the execution of one or more video processing functions by the hardware (e.g., by the graphics module 422). One such function is de-interlacing. Another such function is sub-stream compositing. De-interlacing can be combined with compositing as described above, or either function can be performed separately depending on the received data. For instance, if progressive video data has been received, there is no need to perform de-interlacing; in this case, the VMR module 416 may simply resize the object, add video sub-stream data to it, or perform some other function or combination of functions. Many other functions are be implemented besides de-interlacing and compositing.

Finally, in step 610, the VMR module closes the video stream that it opened in step 606. This step 610 may be in response to a command given by the user, or simply in response to running out of streaming video data, etc.

4 The following discussion presents more fine-grained information regarding
5 selected steps referenced above.

Initialization Steps

8 In step 602 described above, the VMR module 416 asks the display driver 420
9 what processing functionality that it supports with respect to an input video format.
10 When the display driver 420 responds, the VMR module 416 sends a request for more
11 specific information regarding the requirements of the display driver 420 for a particular
12 mode. The display driver 420 responds by specifying various information within a data
13 structure. Such information identifies a number of forward reference samples required, a
14 number of backward reference samples requested, a format of the output frame, etc. A
15 flag is also included in this data structure that indicates whether support for combined de-
16 interlacing and compositing is supported by the graphics module 422 and associated
17 interfaces. This flag is referred to as the DXVA_VideoProcess_Sub-streams flag in an
18 exemplary implementation.

19 Further, in order to correctly support the combined de-interlacing and
20 compositing, the graphics module 422 and associated interfaces and drivers should be
21 capable of independently stretching (horizontally and/or vertically) both the video frame
22 that is being de-interlaced as well as the supplied video sub-streams. This is required in
23 one implementation because the pixel aspect ratio of the primary video and the video sub-
24 streams may be different and non-square in nature. The display driver 420 can
25 communicate its ability to handle this functionality by returning

1 DXVA_VideoProcess_StretchX and DXVA_VideoProcess_StretchY flags that convey
2 its ability to stretch the images, in addition to the DXVA_VideoProcess_Sub-streams
3 flag.

4 Further still, the DeinterlaceBltEx DDI supports significantly enhanced color
5 information for each source and destination surface. The display driver 420 can indicate
6 the level of support it has for this new color information via various color-related flags,
7 e.g.:

8 DXVA_VideoProcess_Sub-streamsExtended;
9 DXVA_VideoProcess_YUV2RGBExtended; and
10 DXVA_VideoProcess_AlphaBlendExtended flags.

11
12 *De-Interlacing Step*

13 The VMR module 416 uses a method referred to as “DeinterlaceBltEx” to
14 coordinate the execution of de-interlacing and compositing by the graphics module 422.
15 More specifically, this DeinterlaceBltEx can be implemented as a single call to the
16 display driver 420, even though it technically involves more than one fundamental
17 operation (de-interlacing and compositing). The DeinterlaceBltEx writes the output of its
18 operation to a specified destination surface.

19 More specifically, the VMR module 416 forwards the following data structure to
20 the display driver 420 to implement the DeinterlaceBltEx method.

21 HRESULT
22 DeinterlaceBltEx(
23 [in] HDXVA_DeinterlaceStream hDiStrm
24 [in] REFERENCE_TIME rtTargetFrame,
25 [in] RECT* prcTargetRect,

```
1     [in] DXVA_AYUVsample2 BackgroundColor,  
2  
3     [in] DWORD DestinationFormat,  
4  
5     [in] DWORD DestinationFlags,  
6  
7     [in] LPDDDSURFACE lpDDSDstSurface,  
8  
9     [in] LPDXVA_VideoSample2 lpDDSSrcSurfaces,  
10  
11    [in] DWORD dwNumSurfaces,  
12  
13    [in] FLOAT Alpha /*0.0F transparent, 1.0F opaque */  
14  
15    );
```

16 In this structure, the rtTargetFrame parameter identifies the temporal location of
17 the output frame within the sequence of input frames. If only de-interlacing is being
18 performed, the target time should coincide with one of the rtStart times or midpoint
19 times, (rtStart+rtEnd)/2, of a reference sample. If a frame rate conversion is being
20 requested, the rtTargetFrame time may be different from any of the rtStart or midpoint
21 times of the reference samples.

22 The prcTargetRect parameter identifies the location within the destination surface
23 that the DeinterlaceBltEx operation should write to. In one implementation, the output
24 should be restricted to the pixels within this rectangle; that is, every pixel within the
25 prcTargetRect should be written to, and pixels outside the prcTargetRect should not be
modified in any way.

26 The BackgroundColor parameter identifies a color of the background upon which
27 all the video stream and sub-streams are composed.

28 The DestinationFormat parameter contains extended color information relating to
29 the destination surface.

30 The DestinationFlags parameter contains a collection of flags that indicate
31 changes in the destination-related parameters from the previous call to DeinterlaceBltEx.

1 These flags reflect changes to the background color, extended color data, target rectangle
2 or the planar alpha parameter and are provided to aid optimizing the driver code.

3 The destination surface will be an off-screen plain surface located in video
4 memory (e.g., local memory 430). The pixel format of the destination surface will be the
5 one indicated in the d3dOutputFormat field of the data structure returned to the VMR
6 module 416 in the initialization steps. In one exemplary implementation, the destination
7 surface specified in the structure should be in the YUV color space.

8 The lpDDSSrcSurfaces parameter points to an array of DXVA_VideoSample2
9 structures. The SampleFormat field in this structure indicates whether the sample is a
10 reference for a de-interlace operation or a video sub-stream sample that needs to be
11 combined with the de-interlaced video frame. Video sub-stream samples should have the
12 DXVA_SampleSub-stream value for their sample format.

13 More specifically, the VideoSample2 structure is identified below.

```
14     typedef struct _DXVA_VideoSample2 {  
15         REFERENCE_TIME     rtStart;  
16         REFERENCE_TIME     rtEnd;  
17         DWORD             SampleFormat;  
18         DWORD             SampleFlags;  
19         VOID*             lpDDSSrcSurface;  
20         RECT              rcSrc;  
21         RECT              rcDst;  
22         DXVA_AYUVsample2  Pal[16];  
23     } DXVA_VideoSample2, *LPDXVA_VideoSample2.  
24  
25
```

1 In addition to indicating whether the sample is interlaced or progressive, the
2 SampleFormat in the above-identified parameter includes extended color information for
3 each sample. This data includes information relating to the: a) color primaries; b) transfer
4 function; c) intended view conditions; d) transfer matrices; and e) black point.

5 For reference video samples, the rtStart and rtEnd fields indicate the temporal
6 location of the sample. For video sub-stream samples, these fields are cleared to zero.

7 The source and destination rectangles are used for subrectangle de-interlacing or
8 stretching. Note that stretching of video sub-streams is independent of the video stream
9 and that support for stretching is mandatory in one implementation. For the palletized
10 video sub-stream pixel formats, a Pal field of the DXVA_VideoSample2 structure
11 contains 16 palette entries that can be used when compositing the sub-stream sample.
12 For non-palletized pixel formats the palette will be cleared to zero and can be ignored.

13 Each input sample contains a collection of flags that indicate changes in the
14 current sample from the previous sample. These flags reflect changes to the palette, color
15 data, source, and destination rectangles of each sample and are provided to aid optimizing
16 the driver code.

17 Continuing with the description of the DeinterlaceBltEx structure, the
18 dwNumSurfaces parameter indicates the number of elements in the lpDDSrcSurface
19 array. The video reference samples will be first in the array followed by the video sub-
20 streams in Z-order. In one exemplary implementation, the number of video sub-streams
21 passed to the driver can range from 0 to 15. Most often when DeinterlaceBltEx is called,
22 the driver will be passed 0 or 1 video sub-streams.

23 Finally, the Alpha parameter indicates a planar transparency value that can be
24 applied to the composite background color, video stream, and sub-stream image as it is
25 written to the destination surface.

1 It should be noted that the DeinterlaceBltEx method can be called by the VMR
2 module 416 when progressive video and multiple video sub-streams are received. This
3 can occur, for instance, when the VMR module 416 is used for DVD playback that
4 contains a mixture of progressive and interlaced video. In this case, the display driver
5 420 will not attempt to de-interlace the video stream (because it is already in the
6 progressive format); rather, the VMR module 416 can be configured to combine the
7 video stream with any given sub-streams, resizing each stream as desired or required. (If
8 a de-interlace mode that needs multiple reference samples is being used with progressive
9 video, the multiple reference samples will still be sent to the display driver 420; however,
10 each reference sample will refer to the same progressive video frame.)

11

12 *Examples of the Operation of DeinterlaceBltEx*

13 The following examples illustrate how the DeinterlaceBltEx method can combine
14 a principal video stream and a video sub-stream (received from, for example, a closed
15 captioning source). Generally, in these examples, the target rectangle (specified in the
16 DeinterlaceBltEx method) identifies the location within the destination surface that the
17 driver 420 should direct its output to. The source and destination rectangles pertain to
18 video stream information specified in the array of VideoSample2 structures (discussed
19 above).

20 The DeinterlaceBltEx method specifies the source rectangle coordinates as
21 absolute locations within the source surface. Likewise, the DeinterlaceBltEx method
22 specifies the destination rectangle coordinates and target rectangle coordinates as
23 absolute locations within the destination surface. In practice, the video stream destination
24 rectangle is often the same size as the destination surface, but this is not always the case.

25 To begin with, Fig. 7 shows an exemplary case 700 in which the DeinterlaceBltEx

1 method simply manipulates an input video frame (without combining it with a video sub-
2 stream). In this example, the input video stream frame 702 has a size of 720x480 pixels
3 and an aspect ratio of 16:9 (where an aspect ratio specifies a ratio of horizontal to vertical
4 dimensions of the video frame). The destination surface has a size of 640x480 pixels and
5 an aspect ratio of 4:3. The target rectangle has coordinates of {0, 0, 640, 480}, the source
6 rectangle has coordinates of {0, 0, 720, 480}, and the destination rectangle has
7 coordinates of {0, 60, 640, 300}.

8 Given these inputs, the DeinterlaceBltEx operation 704 produces the output result
9 706. The output result 706 shows that the video is letter-boxed into the 4:3 destination
10 surface. In performing the DeinterlaceBltEx operation, the driver 420 should de-interlace
11 the video (if necessary), shrink it within the destination surface, and draw the two
12 horizontal bars 708 in the specified background color. This is therefore an example of a
13 case in which the destination rectangle is not the same size as the destination surface.

14 Fig. 8 shows another exemplary case 800 in which the VMR module 416 calls the
15 driver 420 with a video stream destination rectangle that does not fully cover the
16 destination surface. An example of this would be DVD content in which the video
17 stream is specified in the 4:3 aspect ratio and the sub-picture stream is specified in the
18 16:9 aspect ratio. More specifically, as shown in Fig. 8, the input video frame 802 has a
19 size of 720x480 pixels and an aspect ratio of 4:3. The video sub-stream frame 804 has a
20 size of 720x480 pixels and an aspect ratio of 16:9. The destination surface has a size of
21 854x480 pixels and an aspect ratio of 16:9. The target rectangle has coordinates of {0, 0,
22 854, 480}. The source rectangle for the video stream has coordinates of {0, 0, 720, 480}
23 and the corresponding destination rectangle has coordinates of {107, 0, 747, 480}. The
24 source rectangle for the sub-picture stream has coordinates of {0, 0, 720, 480} and the
25 corresponding destination rectangle has coordinates of {0, 0, 854, 480}.

1 Given these inputs, the DeinterlaceBltEx operation 806 produces the output result
2 808. As can be seen, the left and right edges 810 of the destination surface do not contain
3 any pixels from the video stream. When performing the DeinterlaceBltEx operation,
4 pixels that fall outside the video stream's destination sub-rectangle should be interpreted
5 as the background color, as they are combined with the pixels from the sub-picture
6 stream.

7 Fig. 9 shows a case 900 in which the video stream and the video sub-stream have
8 different heights as well as widths. In this case, the input video stream frame 902 has a
9 size of 150x150 pixels and an aspect ratio of 1:1. The video sub-stream frame 904 has a
10 size of 100x85 pixels and an aspect ratio of 1:1. The destination surface has a size of
11 150x85 pixels and an aspect ratio of 1:1. Given these inputs, the DeinterlaceBltEx
12 operation 906 produces the output result 908. As can be seen, the background color is
13 only visible in the four corners 910. The specified background color should only be
14 drawn over the target rectangle.

15 Fig. 10 shows another case 1000 in which the video stream and the video sub-
16 stream have different heights. In this case, the input video frame 1002 has a size of
17 300x150 pixels and an aspect ratio of 1:1. The video sub-stream frame 1004 has a size of
18 150x200 pixels and an aspect ratio of 1:1. The destination surface has a size of 300x200
19 pixels and an aspect ratio of 1:1. The target rectangle has coordinates of {0, 0, 150, 100}.
20 The source rectangle for the video stream has coordinates of {0, 0, 300, 150} and the
21 counterpart destination rectangle has coordinates of {0, 12, 150, 87}. The sub-stream
22 source rectangle has coordinates of {0, 0, 150, 200} and the counterpart destination
23 rectangle has coordinates of {37, 0, 112, 100}.

24 Given these inputs, the DeinterlaceBltEx operation 1006 produces the output
25 result 1008. Here, the VMR module 416 has been instructed to decimate its normal

1 output image size horizontally and vertically by a factor of two. As a result, the
2 background color should only be displayed in the target rectangle, and the remaining
3 pixels 1010 in the destination surface should not be written to. This is because the target
4 rectangle is the bounding rectangle of the video stream and all the sub-streams.

5 Figs. 11-13 show cases in which the rcSrc rectangles (specified in the
6 DXVA_VideoSample2 structure) are different from the source sample sizes. This
7 condition can be indicated by setting a DXVA_VideoProcess_SubRects flag. To begin
8 with, Fig. 11 shows an exemplary case 1100 in which no stretching of video images is
9 performed. In this exemplary case, the destination surface has a size of 720x576 pixels,
10 the rcTarget rectangle has coordinates of {0, 0, 720, 576}, and the background color is
11 solid black. The video stream 1102 is characterized by the following rectangles: the
12 source surface 1104 is demarcated by coordinates of {0, 0, 720, 480}; the rcSrc rectangle
13 1106 is demarcated by coordinates of {360, 240, 720, 480}, and the rcDst rectangle is
14 demarcated by coordinates of {0, 0, 360, 240}. The video sub-stream 1108 is
15 characterized by the following rectangles: the source surface 1110 is demarcated by
16 coordinates of {0, 0, 640, 576}, the rcSrc rectangle 1112 is demarcated by coordinates of
17 {0, 288, 320, 576}, and the rcDst rectangle is demarcated by coordinates of {400, 0, 720,
18 288}.

19 Given these inputs, the DeinterlaceBltEx operation 1114 produces the output
20 result 1116. In this result 1116, the bottom right corner of the video stream (i.e., portion
21 1106) is displayed in the top left corner of the destination surface, and the bottom left
22 corner of the sub-stream (i.e., portion 1112) is displayed in the top right corner of the
23 destination surface. In this example, the hashed areas indicate the sub-rectangles that will
24 be processed.

25 Fig. 12 shows an exemplary case in which the video stream 1202 and the sub-

1 stream 1204 destination rectangles intersect. In this example, the surface dimensions are
2 the same as the example presented above with respect to Fig. 11. The video stream 1202
3 is characterized by a rcSrc rectangle 1206 demarcated by the coordinates of {260, 92,
4 720, 480}, and a rcDst rectangle demarcated by coordinates of {0, 0, 460, 388}. The
5 video sub-stream 1204 is characterized by a rcSrc rectangle 1208 demarcated by
6 coordinates of {0, 0, 460, 388} and a rcDst rectangle demarcated by coordinates of {260,
7 188, 720, 576}. The target rectangle is again demarcated by the coordinates of {0, 0,
8 720, 576}, and the background color is solid black.

9 Given these inputs, the DeinterlaceBltEx operation 1210 produces the output
10 result 1212. In this case, the bottom right corner of the source (i.e., portion 1206) is
11 displayed in the top left corner of the destination, shifted on the X and Y axis by +100.
12 Also, the top left corner of the sub-stream source (i.e., portion 1208) is displayed in the
13 bottom right corner of the destination, shifted on the X and Y axis by -100.

14 Fig. 13 shows a case 1300 where stretching is performed when combining a
15 principal video stream 1302 and a video sub-stream 1304. Namely, in Fig. 13, the
16 destination surface has a size of 720x480 pixels, the target rectangle has a size
17 demarcated by coordinates of {0, 0, 720, 480}, and the background color is black.
18 Further, the video stream source surface 1302 has a size of 360x240 pixels. The video
19 stream 1302 also has source rectangle rcSrc (1306) demarcated by the coordinates of
20 {180, 120, 360, 240} and a destination rectangle rcDst demarcated by the coordinates of
21 {0, 0, 360, 240}. The video sub-stream surface 1308 has a size of 360x240 pixels. The
22 video sub-stream surface 1308 also has a source rectangle rcSrc (1308) demarcated by
23 the coordinates of {0, 0, 180, 120} and a destination rectangle rcDst rectangle
24 demarcated by the coordinates of {360, 240, 720, 480}.

25 Given these inputs, the DeinterlaceBltEx operation 1310 produces the output

1 result 1312. The hashed areas shown in this drawing again indicate the sub-rectangles
2 that are passed to the DeinterlaceBltEx method.

3 *Color Extensions for DeinterlaceBltEx*

4 To improve color fidelity, the apparatus 400 shown in Fig. 4 can also employ a
5 more accurate Y'Cb'Cr' representation when performing color space conversions or
6 video image processing. By way of background, many video formats and standards use
7 different: a) color primaries (e.g., the RGB response functions); b) transfer functions
8 (e.g., the “gamma” function, which governs how to convert R'G'B' to RGB); c) intended
9 viewing conditions (e.g., light or dark rooms); d) transfer matrices (e.g., how to convert
10 Y'Cb'Cr' to R'G'B'; and e) black point (which refers to the conversion to sR'G'B' or to
11 studio R'G'B', including headroom/toeroom).

12 All of the above-described parameters can be encoded into a single 16 bit WORD.
13 In one exemplary implementation, the DeinterlaceBltEx API has a 32 bit DWORD used
14 to describe the format of a sample, such as whether it is interlaced or progressive, and so
15 on. This uses only the lower 8 bits.

16 In the present apparatus 400, the upper 16 bits can be used for conveying extra
17 information used to describe the exact Y'Cb'Cr' colorspace. More specifically, in one
18 exemplary implementation, the 16 bits of information can be packed as follows: a)
19 VideoTransferFunction (4 bits, bits 15-12); b) VideoPrimaries (4 bits, bits 11-8); c)
20 VideoLighting (3 bits, bits 7-5); d) Union, including TransferMatrix (3 bits, 4-2) with
21 respect to the Y'Cb'Cr' image, and NominalRange (3 bits, 4-2) with respect to the RGB
22 image; and e) VideoChromaSubsampling (2 bits, bits 1-0).

23 Each of the above-identified parameters is described below. To begin with, the
24

1 DXVA_VideoTransferFunction enum indicates the conversion function from R'G'B' to
2 RGB. Roughly, it corresponds to the gamma function of the data. Some transfer
3 functions have corrections to account for 8 bit integer quantization effects.

4 The DXVA_VideoPrimaries enum lists the color primaries, stating which RGB
5 basis functions are used.

6 The DXVA_VideoLighting enum describes the intended viewing lighting
7 conditions. These can be used to alter the gamma to generate a comparable experience in
8 a different lighting condition.

9 The DXVA_VideoTransferMatrix enum describes the conversion matrix from
10 Y'Cb'Cr' to (studio) R'G'B'.

11 The DXVA_VideoChromaSubsampling enum describes the chroma encoding
12 scheme for Y'Cb'Cr' data. The 'cosite' variations indicate that the chroma samples are
13 aligned with the luma samples.

14 The DXVA_NominalRange enum describes whether the data includes headroom
15 (values beyond 1.0 white) and toeroom (superblacks below the reference 0.0 black). It is
16 useful to differentiate wide gamut R'G'B' (blackpoint at 16, 16, 16, whitepoint at 235,
17 235, 235) verses normal sRGB.

18 As mentioned above, the following three flags indicate whether the driver 420 is
19 equipped to process certain features of the above-identified extended color functionality:
20

21 DXVA_VideoProcess_SubStreamsExtended;
22 DXVA_VideoProcess_YUV2RGBExtended; and
23 DXVA_VideoProcess_AlphaBlendExtended.

24 Support for the DXVA_VideoProcess_SubStreamsExtended flag indicates that
25 the driver can perform the color adjustments to the source video and sub-streams,

1 indicated in the extended color data parameter, as the video is de-interlaced, composited
2 with the sub-streams, and written to the destination surface. Support for the
3 DXVA_VideoProcess_YUV2RGBExtended flag indicates that the driver is capable of
4 performing a color space conversion operation as the de-interlaced and composited pixels
5 are written to the destination surface using the extended color information that is
6 specified for the source and destination surfaces. Support for the
7 DXVA_VideoProcess_AlphaBlendExtended flag indicates that the driver is capable of
8 performing an alpha blend operation with the destination surface when the de-interlaced
9 and composited pixels are written to the destination surface.

10

11 *Exemplary DDI Mapping for the De-Interlace Interface*

12 According to one exemplary implementation, the DeinterlaceBltEx method maps
13 directly to a RenderMoComp method of the DD_MOTIONCOMPCALLBACKS
14 structure. More specifically, RenderMoComp and DD_MOTIONCOMPCALLBACKS
15 are part of the Device Driver Interface (DDI) between the Microsoft Windows
16 DirectDraw component and the Graphics Device Driver. That is,
17 DD_MOTIONCOMPCALLBACKS is the name of a data structure that contains a table
18 of function pointers and RenderMoComp is one of the functions in the table. When a
19 user calls the DeinterlaceBltEx function it is mapped into a call to the RenderMoComp
20 function in the graphics driver. Different mappings are used for different operating
21 systems. The exemplary mapping described herein is specific to the Windows ®
22 operating system, but the principles described herein can be applied to other operating
23 environments.

24

25 In the mapping: a) dwNumBuffers is 1 + the number of source surfaces; b)

1 lpBufferInfo points to an array of surfaces (where the first surface is the destination
2 surface, and the remaining surfaces are the source surfaces); c) dwFunction is defined as
3 DXVA_DeinterlaceBltExFnCode; lpInputData points to the following exemplary
4 structure:

```
6 #define MAX_DEINTERLACE_INPUT_SURFACES 32
7
8 typedef struct _DXVA_DeinterlaceBltEx {
9
10    DWORD           Size;
11
12    DXVA_AYUVsample2  BackgroundColor;
13
14    RECT            rcTarget;
15
16    REFERENCE_TIME  rtTarget;
17
18    DWORD           NumSourceSurfaces;
19
20    FLOAT           Alpha;
21
22    DXVA_VideoSample2  Source[MAX_DEINTERLACE_SURFACES];
23
24    DWORD           DestinationFormat;
25
26    DWORD           DestinationFlags;
27
28 } DXVA_DeinterlaceBltEx;
```

17 and d) lpOutputData is NULL. In one exemplary implementation, for the DX-VA device
18 used for de-interlacing, RenderMoComp will be called without calling
19 BeginMoCompFrame or EndMoCompFrame.

20 In the above structure, the `lpBufferInfo` parameter points to an array of surfaces.
21 In one exemplary implementation, the order of the surfaces within the array obeys the
22 following rules. The first surface in the array is the Destination surface; this is the only
23 surface in the array that should be written too. The next sequence of surfaces in the array
24 would be any previous destination surfaces, in reverse temporal order, that the de-

1 interlacing device requested for it's de-interlace algorithm. The next sequence of
2 surfaces in the array is a collection of input interlaced or progressive surfaces that the
3 device needs in order to perform it's de-interlace operation. The next sequence of
4 surfaces in the array is the video sub-stream surfaces, these surfaces being in Z order.

5 The following exemplary sample code below shows how the driver can map the
6 RenderMoComp DDI call into calls to DeinterlaceBltEx. The sample code only shows
7 how the RenderMoComp function is used for de-interlacing. If the driver supports other
8 DX-VA functions, such as decoding MPEG-2 video streams, then the sample code can be
9 extended to include processing of additional DX-VA GUIDs.

```
10
11    DWORD APIENTRY
12
13        RenderMoComp (
14
15            PDD_RENDERMOCOMPDATA lpData
16
17            )
18
19        {
20
21            LPDXVA_DeinterlaceStream pDXVAState =
22
23                (LPDXVA_DeinterlaceStream)lpData->lpMoComp->lpDriverReserved1;
24
25            DXVA_DeinterlaceBltEx* lpBlt =
26
27                (DXVA_DeinterlaceBltEx*)lpData->lpInputData;
28
29            LPDDMOCOMPBUFFERINFO lpBuffInfo = lpData->BufferInfo;
30
31            for (DWORD i = 0; i < lpBlt->NumSourceSurfaces; i++) {
32
33                lpBlt->Source[i].lpDDSSrcSurface =
34
35                    lpBuffInfo[1 + i].lpCompSurface;
36
37            }
38
39            lpData->ddRVal = DeinterlaceBltEx(pDXVAState,
40
41
42                lpBlt->rtTarget,
43
44                &lpBlt->rcTarget,
```

```
1           lpBlt->BackgroundColor,
2
3           lpBlt->DestinationFormat,
4
5           lpBlt->DestinationFlags,
6
7           lpBuffInfo[0].lpCompSurface,
8
9           &lpBlt->Source,
10
11          lpBlt->NumSourceSurfaces,
12
13          lpBlt->Alpha);
14
15
16      return DDHAL_DRIVER_HANDLED;
17
18  }
```

9 **D. Exemplary Computer Environment**

10 Fig. 14 illustrates one example of a computing environment 1400 within which
11 the above-described video processing can be either fully or partially implemented. The
12 computing environment 1400 includes the general purpose computer 1402 and the
13 display device 432 (discussed in the context of Fig. 4). However, the computing
14 environment 1400 can include other kinds of computer and network architectures. For
15 example, although not shown, the computer environment 1400 can include hand-held or
16 laptop devices, set top boxes, programmable consumer electronics, mainframe
17 computers, gaming consoles, etc. Further, Fig. 14 shows elements of the computer
18 environment 1400 grouped together to facilitate discussion. However, the computing
19 environment 1400 can employ a distributed processing configuration. In a distributed
20 computing environment, computing resources can be physically dispersed throughout the
21 environment.

22 Exemplary computer 1402 includes one or more processors or processing units
23 1404, a system memory 1406, and a bus 1408. The bus 1408 connects various system
24 components together. For instance, the bus 1408 connects the processor 1404 to the
25 system memory 1406. The bus 1408 can be implemented using any kind of bus structure

1 or combination of bus structures, including a memory bus or memory controller, a
2 peripheral bus, an accelerated graphics port, and a processor or local bus using any of a
3 variety of bus architectures. For example, such architectures can include an Industry
4 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced
5 ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a
6 Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

7 Computer 1402 can also include a variety of computer readable media, including
8 a variety of types of volatile and non-volatile media, each of which can be removable or
9 non-removable. For example, system memory 1406 includes computer readable media in
10 the form of volatile memory, such as random access memory (RAM) 1410, and non-
11 volatile memory, such as read only memory (ROM) 1412. ROM 1412 includes an
12 input/output system (BIOS) 1414 that contains the basic routines that help to transfer
13 information between elements within computer 1402, such as during start-up. RAM
14 1410 typically contains data and/or program modules in a form that can be quickly
15 accessed by processing unit 1404.

16 Other kinds of computer storage media include a hard disk drive 1416 for reading
17 from and writing to a non-removable, non-volatile magnetic media, a magnetic disk drive
18 1418 for reading from and writing to a removable, non-volatile magnetic disk 1420 (e.g.,
19 a “floppy disk”), and an optical disk drive 1422 for reading from and/or writing to a
20 removable, non-volatile optical disk 1424 such as a CD-ROM, DVD-ROM, or other
21 optical media. The hard disk drive 1416, magnetic disk drive 1418, and optical disk drive
22 1422 are each connected to the system bus 1408 by one or more data media interfaces
23 1426. Alternatively, the hard disk drive 1416, magnetic disk drive 1418, and optical disk
24 drive 1422 can be connected to the system bus 1408 by a SCSI interface (not shown), or
25 other coupling mechanism. Although not shown, the computer 1402 can include other

1 types of computer readable media, such as magnetic cassettes or other magnetic storage
2 devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical
3 storage, electrically erasable programmable read-only memory (EEPROM), etc.

4 Generally, the above-identified computer readable media provide non-volatile
5 storage of computer readable instructions, data structures, program modules, and other
6 data for use by computer 1402. For instance, the readable media can store the operating
7 system 1428, one or more application programs 1430, other program modules 1432, and
8 program data 1434.

9 The computer environment 1400 can include a variety of input devices. For
10 instance, the computer environment 1400 includes the keyboard 1436 and a pointing
11 device 1438 (e.g., a "mouse") for entering commands and information into computer
12 1402. The computer environment 1400 can include other input devices (not illustrated),
13 such as a microphone, joystick, game pad, satellite dish, serial port, scanner, card reading
14 devices, digital or video camera, etc. Input/output interfaces 1439 couple the input
15 devices to the processing unit 1404. More generally, input devices can be coupled to the
16 computer 1402 through any kind of interface and bus structures, such as a parallel port,
17 serial port, game port, universal serial bus (USB) port, etc.

18 The computer environment 1400 also includes the display device 432.
19 Functionality 1440 generally represents the vertically disposed elements shown in Fig. 4,
20 for example, renderer 416, graphics interface 418, display driver 420, graphics processor
21 422, and so on. In addition to the display device 432, the computer environment 1400
22 can include other output peripheral devices, such as speakers (not shown), a printer (not
23 shown), etc.

24 Computer 1402 can operate in a networked environment using logical connections
25 to one or more remote computers, such as a remote computing device 1442. The remote

1 computing device 1442 can comprise any kind of computer equipment, including a
2 general purpose personal computer, portable computer, a server, a router, a network
3 computer, a peer device or other common network node, etc. Remote computing device
4 1442 can include all of the features discussed above with respect to computer 1402, or
5 some subset thereof.

6 Any type of network can be used to couple the computer 1402 with remote
7 computing device 1442, such as a local area network (LAN) 1444, or a wide area
8 network (WAN) 1446 (such as the Internet). When implemented in a LAN networking
9 environment, the computer 1402 connects to local network 1444 via a network interface
10 or adapter 1448. When implemented in a WAN networking environment, the computer
11 1402 can connect to the WAN 1446 via a modem 1450 or other connection strategy. The
12 modem 1450 can be located internal or external to computer 1402, and can be connected
13 to the bus 1408 via serial I/O interfaces 1452 other appropriate coupling mechanism.
14 Although not illustrated, the computing environment 1400 can provide wireless
15 communication functionality for connecting computer 1402 with remote computing
16 device 1442 (e.g., via modulated radio signals, modulated infrared signals, etc.).

17 In a networked environment, the computer 1402 can draw from program modules
18 stored in a remote memory storage device 1454. Generally, the depiction of program
19 modules as discrete blocks in Fig. 14 serves only to facilitate discussion; in actuality, the
20 programs modules can be distributed over the computing environment 1400, and this
21 distribution can change in a dynamic fashion as the modules are executed by the
22 processing unit 1404.

23
24 In closing, although the invention has been described in language specific to
25 structural features and/or methodological acts, it is to be understood that the invention

1 defined in the appended claims is not necessarily limited to the specific features or acts
2 described. Rather, the specific features and acts are disclosed as exemplary forms of
3 implementing the claimed invention.

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25